

We claim:

1. A method for constructing and caching a chain of file identifiers that represent a full path to a file system resource comprising the steps of:

retrieving a file identifier corresponding to the file system resource which is the
5 target of the access attempt and a file identifier chain for the directory of the target system resource;

searching for the effective security classification category and defined name for the target resource file identifier;

updating the security classification system, when said search finds a security
10 classification category for the target resource file identifier;

determining whether operations for the target file system resource could affect the file system name space; and

terminating said method when operation does not affect the file system name space.
15

2. The method as described in claim 1 wherein after said searching step, the security classification category is set to an unclassified category and the defined name is set to the path used in the file system resource access attempt when said security classification category search does not find a security classification category.
20

3. The method as described in claim 1 further comprising the step of flushing the a file identifier chain cache when there is a determination that desired operations on the target file system resource could affect the file system name space.

25 4. The method as described in claim 1 further comprising before said file identifier (FID) retrieval step the step of processing a system resources defined name (DN) and security classification category into a mapping database which holds a FID to DN mapping.

5. The method as described in claim 4 wherein said database processing step comprises:

providing the defined name and security classification category as inputs;

obtaining a file identifier (FID) for the defined name; and

5 adding the FID to DN mapping containing the security classification category to the mapping database.

6. The method as described in claim 1 wherein said searching step comprises:

searching the FID to DN mapping database for the security classification category

10 for the FID of the target resource; and

returning the security classification category and defined name for the target FID, when a security classification category for the target FID was found during said search.

7. The method as described in claim 1 wherein said searching step comprises:

15 searching the FID to DN mapping database for the security classification category for the FID of the target resource;

retrieving a FID from the FID chain, when the search does not find a security classification category for the FID of the target resource;

20 searching the FID to DN mapping database for the security classification category for the FID of the FID chain; and

returning the security classification category and defined name for the target FID, when a security classification category for the target FID was found during said search.

8. The method as described in claim 7 further comprising the steps of:

25 determining whether more entries in the FID chain, when the search does not find a security classification category for the FID used in the search;

retrieving the next FID in the FID chain; and

searching the FID to DN mapping database for the security classification category for the currently retrieved FID of the FID chain.

9. The method as described in claim 8 further comprising the step of terminating the method when no security classification category is found for any FID in the FID chain.

10. The method as described in claim 3 wherein said flushing step comprises:

5 retrieving the path name for the target resource, said path name being to a directory for the target resource;

obtaining a vnode for the directory;

generating a FID for the directory using the vnode;

searching for FID chain matching directory FID; and

10 removing FID chain from cache, when matching FID chain is found.

11. The method as described in claim 10 further comprising before said searching step the step of sorting the FID chains in the FID chain cache into hash list.

15 12. The method as described in claim 11 wherein said searching step comprises:

retrieving the first FID chain in the FID chain list;

comparing each FID in said first FID chain to said directory FID;

determining whether there are more FID chains in the list, when said FID chain did not match said directory FID;

20 retrieving the next FID chain in the FID, and

returning to said comparing step using newly retrieved FID chain.

13. The method as described in claim 11 wherein said searching step comprises: retrieving the first FID chain in the FID chain list;

25 comparing each FID in said first FID chain to said directory FID;

determining whether there are more FID chains in the list, when said FID chain did not match said directory FID; and

terminating method when no FID chain is found.

14. A computer program product in a computer readable medium for use in constructing and caching a chain of file identifiers that represent a full path to a file system resource comprising:

instructions for retrieving a file identifier corresponding to the file system resource which is the target of the access attempt and a file identifier chain for the directory of the target system resource;

instructions for searching for the effective security classification category and defined name for the target resource file identifier;

instructions for updating the security classification system, when said search finds a security classification category for the target resource file identifier;

instructions for determining whether operations for the target file system resource could affect the file system name space; and

instructions for terminating said method when operation does not affect the file system name space.

15. The computer program product as described in claim 14 further comprising instructions for flushing the a file identifier chain cache when there is a determination that desired operations on the target file system resource could affect the file system name space.

16. The computer program product as described in claim 15 wherein said flushing instructions comprise:

instructions for retrieving the path name for the target resource, said path name being to a directory for the target resource;

instructions for obtaining a vnode for the directory;

instructions for generating a FID for the directory using the vnode;

instructions for searching for FID chain matching directory FID; and

instructions for removing FID chain from cache, when matching FID chain is found.

17. The computer program product as described in claim 14 wherein said searching instruction comprises:

instructions for searching the FID to DN mapping database for the security classification category for the FID of the target resource;

5 instructions for retrieving a FID from the FID chain, when the search does not find a security classification category for the FID of the target resource;

instructions for searching the FID to DN mapping database for the security classification category for the FID of the FID chain; and

10 instructions for returning the security classification category and defined name for the target FID, when a security classification category for the target FID was found during said search.

18. The computer program product as described in claim 17 further comprising the steps of:

15 instructions for determining whether more entries in the FID chain, when the search does not find a security classification category for the FID used in the search;

instructions for retrieving the next FID in the FID chain; and

instructions for searching the FID to DN mapping database for the security classification category for the currently retrieved FID of the FID chain.

19. The computer program product as described in claim 18 further comprising before said searching instructions, instructions for sorting the FID chains in the FID chain cache into hash list.

20. The computer program product as described in claim 19 wherein said searching instruction comprises:

instructions for retrieving the first FID chain in the FID chain list;

instructions for comparing each FID in said first FID chain to said directory FID;

instructions for determining whether there are more FID chains in the list, when

30 said FID chain did not match said directory FID; and

instructions for terminating method when no FID chain is found.

21. The method as described in claim 1 wherein said file identifier retrieval step comprises:

retrieving the path name of the file resource which is the target of the access attempt;

5 obtaining a FID for target resource with said path name;
determining whether obtained FID is in a FID chain; and
returning the target FID and FID chain, when the target resource FID was found in the FID Chain Cache.

10 22. The method as described in claim 21 further comprising after said path name retrieval step, the step of obtaining vnodes for the target path and parent directory.

23. The method as described in claim 1 wherein said file identifier retrieval step comprises:

15 retrieving the path name of the file resource which is the target of the access attempt;

obtaining a FID for target resource with said path name;
determining whether obtained FID is in a FID chain; and
constructing a FID chain for the parent directory, when no FID chain is found.

20

24. The method as described in claim 23 wherein said FID chain construction comprises:

setting a temporary vnode to equal the vnode for the parent of the target resource;
determining whether the temporary vnode is the root directory;

25 inserting FID chain into FID chain cache with the first FID in the chain serving as the entry search key, when temporary vnode is the root directory.

25. The method as described in claim 23 wherein said FID chain construction comprises:

setting a temporary vnode to equal the vnode for the parent of the target resource;

- 5 determining whether the temporary vnode is the root directory;
retrieving a vnode for the next parent in the directory path and determining whether that parent is the root directory;
repeating said retrieving step until parent is the root of the directory.

10 26. The method as described in claim 25 further comprising the step of inserting a completed FID chain into the FID chain cache when the parent is the root directory.

27. A computer connectable to a distributed computing system which includes file system objects containing information accessed during the execution of application and

- 15 system programs comprising:
a processor;
a native operating system;
application programs;
an external authorization program overlaying said native operating system and
20 augmenting standard security controls of said native operating system;
a cache storage location for store file identifier chains which represent paths to system resources, said cache providing for faster searches of file identifiers.
an access decision component within said external authorization program for determining access to protected file system objects.

25 28. The method as described in claim 1 wherein said method is implemented through the use of externally stored attributes, said attributes being security rules for system resources and further comprising the step of attaching security rules of a directory to all files in said directory.